

What is claimed is:

- 1 1. A method of communicating between a source and a target node in a clustered
2 computer system, the method comprising:
 - 3 establishing a cluster data port between the source node and a target node,
4 the cluster data port configured to select among a plurality of connection paths
5 between the source node and the target node, and to selectively switch over data
6 flow from the target node to a backup target node; and
7 communicating data from the source node to the target node using the
8 cluster data port.

- 1 2. The method of claim 1, wherein establishing the cluster data port includes
2 establishing a plurality of logical connections between the source node and the target
3 node, each logical connection configured to communicate data over a connection path
4 among the plurality of connection paths.

- 1 3. The method of claim 2, wherein establishing the cluster data port includes
2 exchanging between the source and target nodes network addresses associated with at
3 least one of the source and target nodes.

- 1 4. The method of claim 2, wherein establishing the cluster data port includes
2 registering a client that is resident on the source node with the cluster data port.

- 1 5. The method of claim 4, wherein registering the client comprises identifying to
2 the data cluster port at least one callback function associated with the client, wherein the
3 cluster data port is configured to notify the client of a data port event by calling the
4 callback function.

1 6. The method of claim 1, wherein communicating data from the source node to
2 the target node includes performing load balancing in the cluster data port to distribute
3 the data among the plurality of connection paths.

1 7. The method of claim 1, wherein each connection path comprises a TCP
2 connection.

1 8. The method of claim 1, wherein the cluster data port is configured to
2 communicate data between the source and target nodes according to a messaging protocol
3 selected from the group consisting of an asynchronous messaging protocol and a
4 synchronous messaging protocol.

1 9. The method of claim 1, further comprising, with the cluster data port,
2 switching data flow from the target node to a backup target node in response to an
3 inability to communicate with the target node.

1 10. The method of claim 9, wherein switching data flow from the target node to
2 the backup target node includes establishing a logical connection between the source
3 node and the backup target node.

1 11. The method of claim 9, further comprising notifying a client of the cluster
2 data port service of the inability to communicate with the target node, wherein switching
3 data flow from the target node to a backup node is performed in response to initiation of a
4 cluster data port failover by the client.

1 12. The method of claim 9, wherein switching data flow from the target node to a
2 backup node is initiated by the cluster data port.

1 13. The method of claim 1, further comprising, with the cluster data port,
2 resetting a logical connection between the source node and the target node in response to
3 an inability to communicate with the target node.

1 14. The method of claim 1, further comprising communicating data from the
2 target node to the source node using the cluster data port.

1 15. The method of claim 1, wherein the target node is remote from the source
2 node, and wherein communicating data from the source node to the target node includes
3 communicating mirror data to support remote mirroring between the source and target
4 nodes.

- 1 16. An apparatus, comprising:
 - 2 a memory;
 - 3 at least one processor; and
 - 4 program code resident in the memory and configured for execution on the
5 at least one processor to implement a cluster data port for a clustered computer
6 system of the type including a plurality of nodes, the cluster data port configured
7 to support communication between a source node and a target node among the
8 plurality of nodes, the cluster data port further configured to select among a
9 plurality of connection paths between the source node and the target node, and to
10 selectively switch over data flow from the target node to a backup target node.
- 1 17. The apparatus of claim 16, wherein the cluster data port is configured to
2 establish a plurality of logical connections between the source node and the target node,
3 each logical connection configured to communicate data over a connection path among
4 the plurality of connection paths.
- 1 18. The apparatus of claim 17, wherein the cluster data port is configured to
2 exchange between the source and target nodes network addresses associated with at least
3 one of the source and target nodes.
- 1 19. The apparatus of claim 17, wherein the cluster data port is configured to
2 register a client that is resident on the source node.
- 1 20. The apparatus of claim 19, wherein the cluster data port is configured to
2 receive from the client at least one callback function associated with the client during
3 registration of the client, wherein the cluster data port is configured to notify the client of
4 a data port event by calling the callback function.

1 21. The apparatus of claim 16, wherein the cluster data port is configured to load
2 balance data communicated between the source and target nodes to distribute the data
3 among the plurality of connection paths.

1 22. The apparatus of claim 16, wherein each connection path comprises a TCP
2 connection.

1 23. The apparatus of claim 16, wherein the cluster data port is configured to
2 communicate data between the source and target nodes according to a messaging protocol
3 selected from the group consisting of an asynchronous messaging protocol and a
4 synchronous messaging protocol.

1 24. The apparatus of claim 16, wherein the cluster data port is configured to
2 switch data flow from the target node to a backup target node in response to an inability
3 to communicate with the target node.

1 25. The apparatus of claim 24, wherein the cluster data port is configured to
2 switch data flow from the target node to the backup target node by establishing a logical
3 connection between the source node and the backup target node.

1 26. The apparatus of claim 24, wherein the cluster data port is further configured
2 to notify a client of the cluster data port of the inability to communicate with the target
3 node, and wherein the cluster data port is configured to switch data flow from the target
4 node to a backup node in response to initiation of a cluster data port failover by the client.

1 27. The apparatus of claim 24, wherein the cluster data port is configured to
2 initiate the switch of data flow from the target node to a backup node.

1 28. The apparatus of claim 16, wherein the cluster data port is configured to reset
2 a logical connection between the source node and the target node in response to an
3 inability to communicate with the target node.

1 29. The apparatus of claim 16, wherein the cluster data port is configured to
2 support bidirectional communication between the source and target nodes.

1 30. The apparatus of claim 16, wherein the target node is remote from the source
2 node, and wherein the cluster data port is configured to communicate mirror data from the
3 source node to the target node to support remote mirroring between the source and target
4 nodes.

- 1 31. A clustered computer system, comprising:
 - 2 a plurality of nodes; and
 - 3 a cluster data port resident on at least one of the plurality of nodes and
 - 4 configured to support communication between a source node and a target node
 - 5 among the plurality of nodes, the cluster data port configured to select among a
 - 6 plurality of connection paths between the source node and the target node, and to
 - 7 selectively switch over data flow from the target node to a backup target node.
- 1 32. The clustered computer system of claim 31, wherein the cluster data port is
- 2 configured to select among the plurality of connection paths using a load balancing
- 3 algorithm.
- 1 33. The clustered computer system of claim 31, wherein the cluster data port is
- 2 configured to switch over data flow from the target node to the backup target node in
- 3 response to an inability of the source node to communicate with the target node.
- 1 34. The clustered computer system of claim 31, wherein the target node is remote
- 2 from the source node, and wherein the cluster data port is configured to communicate
- 3 mirror data from the source node to the target node to support remote mirroring between
- 4 the source and target nodes.

1 35. A program product, comprising:

2 program code configured to implement a cluster data port for a clustered
3 computer system of the type including a plurality of nodes, the cluster data port
4 configured to support communication between a source node and a target node
5 among the plurality of nodes, the cluster data port further configured to select
6 among a plurality of connection paths between the source node and the target
7 node, and to selectively switch over data flow from the target node to a backup
8 target node; and

9 a signal bearing medium bearing the program code.

1 36. The program product of claim 35, wherein the signal bearing medium
2 includes at least one of a recordable and a transmission medium.